

(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(2)

(11)

EP 1 041 774 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
04.10.2000 Bulletin 2000/40(51) Int. Cl.⁷: H04L 12/56, H04Q 11/04

(21) Application number: 00302447.8

(22) Date of filing: 24.03.2000

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 31.03.1999 US 282980

(71) Applicant: Fore Systems, Inc.
Warrendale, Pennsylvania 15086 (US)

(72) Inventors:
• Kim, Hyong S.
Pittsburgh, PA 15215 (US)
• Basak, Debasish
Allison Park, PA 15101 (US)

(74) Representative:
O'Connell, David Christopher
Haseltine Lake & Co.,
Imperial House,
15-19 Kingsway
London WC2B 6UD (GB)

(54) Method and apparatus for dynamic call admission control

(57) A switch for a network. The switch includes an input port mechanism which receives traffic of connections from the network. Each connection has a priority. The switch includes an output port mechanism which sends traffic of connections to the network. The switch includes a controller which serves connections and which monitors the connections received by the input port mechanism and sent by the output port mechanism and releases connections according to a connection's priority when a predetermined condition exists in the

switch. The controller is connected to the input port mechanism and the output port mechanism. Each connection requests a specific bandwidth from the controller. A method for switching connections. The method includes the steps of monitoring traffic of connections received by a switch. Then there is the step of releasing connections from the switch according to the connection's priority when a predetermined condition in the switch exists.

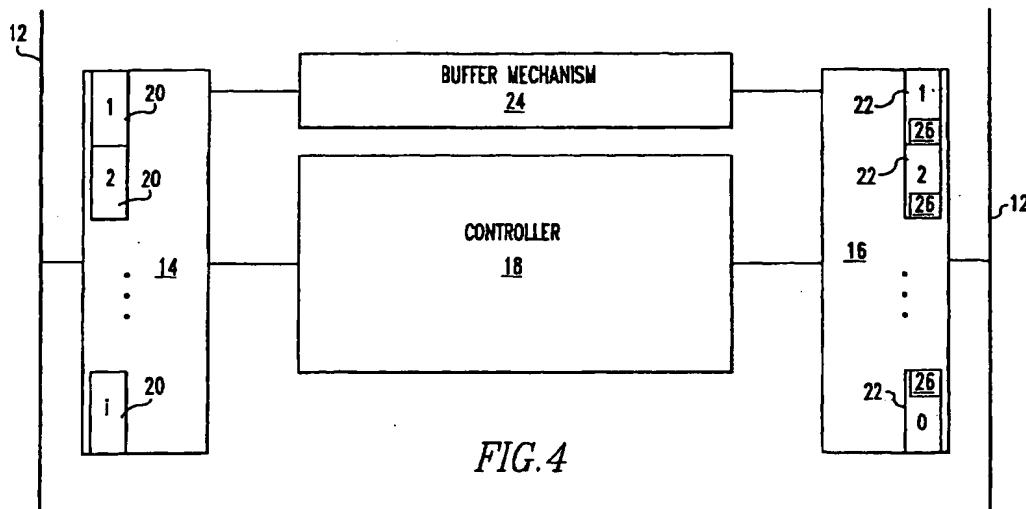


FIG. 4

Description**FIELD OF THE INVENTION**

5 [0001] The present invention pertains to a switch which releases connections according to a connection's priority when a predetermined condition exists in the switch. More specifically, the present invention pertains to a switch which releases connections according to a connection's priority when an overload condition exists in the switch.

BACKGROUND OF THE INVENTION

10 [0002] Current call admission control (CAC) algorithms generally are based only on traffic characteristics specified by a user. However, for most connections, it is very difficult to predict the nature of traffic accurately. Typically, the traffic characteristics of such connections are either over- or under-estimated. Let us consider the impact on overall link utilization while relying on traffic characteristics specified by user. Under-estimation of traffic can lead to excess traffic to link. However, with policing/shaping of traffic at the user network interface (UNI), this situation is avoided. Over-estimation, on the other hand, can lead to the resultant traffic being lower than expected leading to under-utilization of link bandwidth.

15 [0003] High-speed fast-packet-switched network architectures are capable of supporting a wide range of connections with different bandwidth requirements and traffic characteristics. Providing quality of service (QoS) in networks in terms of delay and bandwidth requires provisioning of network resources. While this environment provides increased flexibility in supporting various services, its dynamic nature poses difficult traffic control problems when trying to achieve efficient use of network resources. One such problem is the issue of bandwidth management and allocation. Due to the statistical multiplexing of all connections at the physical layer and the varying transmission rates by connections, it is important to characterize, both the effective bandwidth of a single connection and several multiplexed connections. 20 Such characterization can be used to compute metrics for efficient bandwidth management, routing, and call control procedures. Based on statistical characteristics and desired grade of service, researchers in R. Guerin, H. Ahmadi, and M. Naghshineh. Equivalent capacity and its application to bandwidth allocation in high-speed networks. *IEEE Journal on Selected Areas in Communication*, 9(7) :968-981, Sept. 1991, incorporated by reference herein, have proposed approximate expressions for effective bandwidth or equivalent capacity computations.

25 [0004] A CAC algorithm based on a metric of effective bandwidth, can be summarized as follows. At any given time a new connection is admitted if:

$$B_{\text{eff}}^c + B_{\text{eff}}' \leq B_{\text{link}} \quad (1)$$

30 where, B_{eff}^c denotes the effective bandwidth of new connection, B_{eff}' denotes the effective bandwidth of aggregate traffic on a link, and B_{link} denotes the maximum link capacity or *line rate*. If the connection is admitted, the effective bandwidth of link is updated as, $B_{\text{eff}}' + B_{\text{eff}}^c = B_{\text{eff}}'$, otherwise it remains unchanged. Similarly, a tear-down of an established connection results in link effective bandwidth being decremented by the effective bandwidth of torn-down connection.

35 [0005] While the above approach is simple, it relies on the user/network being able to predetermine the traffic characteristics of a connection. However, for most connections, it is very difficult to predict the nature of traffic accurately. Typically, the traffic characteristics of such connections are either over- or under-estimated. Let us consider the impact on overall link utilization while relying on traffic characteristics specified by user. Under-estimation of traffic can lead to excess traffic to link. However, with policing/shaping of traffic at the user network interface (UNI), this situation is avoided. Over-estimation, on the other hand, can lead to the resultant traffic being lower than expected leading to under-utilization of link bandwidth.

40 [0006] The present invention provides a dynamic CAC to overcome the problem of under-utilization of link bandwidth. The scheme, using periodic measurements, keeps track of various utilization parameters. These include the current link utilization, the variation in link utilization, buffer occupancy, and rate of change of buffer occupancy. Based on 45 the above measurements a dynamic estimate of the effective bandwidth of the aggregate traffic on link is computed and admittance of new calls judged based on this.

SUMMARY OF THE INVENTION

50 [0007] The present invention pertains to a switch for a network. The switch comprises an input port mechanism which receives traffic of connections from the network. Each connection has a priority. The switch comprises an output port mechanism which sends traffic of connections to the network. The switch comprises a controller which serves connections and which monitors the connections received by the input port mechanism and sent by the output port mech-

anism and releases connections according to a connection's priority when a predetermined condition exists in the switch. The controller is connected to the input port mechanism and the output port mechanism. Each connection requests a specific bandwidth from the controller.

[0008] The present invention pertains to a method for switching connections. The method comprises the steps of monitoring traffic of connections received by a switch. Then there is the step of releasing connections from the switch according to the connection's priority when a predetermined condition in the switch exists.

[0009] The present invention pertains to a switch for a network. The switch comprises an input port mechanism which receives traffic of connections from the network. The switch comprises an output port mechanism which sends traffic of connections to the network. The switch comprises a controller which provides service to traffic of connections received at the input port mechanism and which are to be sent out the output port mechanism. The controller monitoring the traffic of connections received by the input port mechanism and sent out the output port mechanism and adjusting the service provided to the connections based on the traffic of connections received by the input port mechanism and sent out the output port mechanism. The controller is connected to the input port mechanism and the output port mechanism.

[0010] The present invention pertains to a switch for a network. The switch comprises an input port mechanism which receives cells of connections from the network. The switch comprises an output port mechanism which sends cells of connections to the network. The switch comprises a buffer mechanism for storing cells. The buffer mechanism is connected to the input port mechanism and the output port mechanism. The switch comprises a controller which provides service to traffic of connections received at the input port mechanism and which are to be sent out the output port mechanism. The controller monitors the change in the number of cells in the buffer mechanism and adjusts the service provided to the connections based on the change in the number of cells in the buffer mechanism. The controller is connected to the buffer mechanism.

[0011] The present invention pertains to a method for switching connections by a switch of a network. The method comprises the steps of monitoring the traffic of connections received by the switch and sent out the switch. Then there is the step of adjusting the service provided to the connections by the switch based on the traffic of connections received by the switch and sent out the switch.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] In the accompanying drawings, the preferred embodiment of the invention and preferred methods of practicing the invention are illustrated in which:

Figures 1a-1g show characteristics of calls generated in simulations.

Figures 2a-2e show plots which depict observed number of calls admitted by gcac, number of calls rejected, gcac parameters, link utilization, and vbr buffer occupancy.

Figures 3a-3f are plots which depict observed number of calls admitted by dcac, number of calls rejected, dcac parameters, d-gcac parameters, link utilization, and vbr buffer occupancy. Note average link utilization with DCAC, remaining the pre-set 80% line rate or 292,400 cells/s, is much higher with dcac than with gcac as shown in earlier plots.

Figure 4 is a schematic representation of a switch of the present invention.

DETAILED DESCRIPTION

[0013] Referring now to the drawings wherein like reference numerals refer to similar or identical parts throughout the several views, and more specifically to figure 4 thereof, there is shown a switch 10 for a network 12. The switch 10 comprises an input port mechanism 14 which receives traffic of connections from the network 12. Each connection has a priority. The switch 10 comprises an output port mechanism 16 which sends traffic of connections to the network 12. The switch 10 comprises a controller 18 which serves connections and which monitors the connections received by the input port mechanism 14 and sent by the output port mechanism 16 and releases connections according to a connection's priority when a predetermined condition exists in the switch 10. The controller 18 is connected to the input port mechanism 14 and the output port mechanism 16. Each connection requests a specific bandwidth from the controller 18.

[0014] Preferably, the predetermined condition includes an overload condition of the switch 10. The overload condition preferably includes receiving more traffic on connections received by the input port mechanism 14 than the controller 18 can process at a given time. Preferably, the controller 18 determines excess bandwidth associated with the

overload condition, and the controller 18 releases the connection, for a given priority, with the closest match to the excess bandwidth.

[0015] The controller 18 preferably releases connections according to the connections' traffic class, where UBR connections are released first, then ABR connections are released next, then VBR connections are released next and then CBR connections are released from the switch 10. Preferably, the input port mechanism 14 comprises I input ports 20 connected to the controller 18, where I is an integer and greater than or equal to 1; and wherein the output mechanism comprises O output ports 22 connected to the controller 18, where O is an integer greater than or equal to 1. The controller 18 preferably releases the connection with the lowest priority.

[0016] Preferably, the controller 18 releases connections whose sum of requested bandwidths most closely matches the excess bandwidth. Preferably, the switch includes a buffer mechanism connected to the output port mechanism. The overload condition preferably exists if the sum of requested bandwidths of connections on an output port is greater than the output port's bandwidth and buffer resources associated with it are being exceeded

$$\text{if } (Q_{vcbr} > T) \text{ and } (B'_{\text{eff}} > B_{\text{link}})$$

where Q_{vcbr} is the buffer occupancy for VBR/CBR traffic, T is a threshold, B'_{eff} is the aggregate effective bandwidth of connections on link 1, and B_{link} is the maximum link capacity or line rate.

[0017] The present invention pertains to a method for switch 10ing connections. The method comprises the steps of monitoring traffic of connections received by a switch 10. Then there is the step of releasing connections from the switch 10 according to the connection's priority when a predetermined condition in the switch 10 exists.

[0018] Preferably, the monitoring step includes the step of monitoring traffic of connections received by and sent out from the switch 10. The releasing step preferably includes the step of releasing connections when the predetermined condition of overload exists in the switch 10 where there is more traffic of connections received by the switch 10 than the switch 10 can process.

[0019] Preferably, the releasing step includes the steps of determining, when the overload condition exists, the excess bandwidth associated with the overload condition; choosing a connection, for a given priority, with the closest match to the excess bandwidth; and releasing the connection with the closest match to the excess bandwidth. The choosing step preferably includes the step of choosing the connection with the lowest priority.

[0020] Preferably, the choosing step includes the step of choosing the connections whose sum of requested bandwidths most closely matches the excess bandwidth. The choosing step preferably includes the step of choosing a connection according to traffic class for release, where UBR connections are released first, then ABR connections are released next, then VBR connections are released next and then CBR connections are released.

[0021] The present invention pertains to a switch 10 for a network 12. The switch 10 comprises an input port mechanism 14 which receives traffic of connections from the network 12. The switch 10 comprises an output port mechanism 16 which sends traffic of connections to the network 12. The switch 10 comprises a controller 18 which provides service to traffic of connections received at the input port mechanism 14 and which are to be sent out the output port mechanism 16. The controller 18 monitors the traffic of connections received by the input port mechanism 14 and sent out the output port mechanism 16 and adjusting the service provided to the connections based on the traffic of connections received by the input port mechanism 14 and sent out the output port mechanism 16. The controller 18 is connected to the input port mechanism 14 and the output port mechanism 16.

[0022] The present invention pertains to a switch 10 for a network 12. The switch 10 comprises an input port mechanism 14 which receives cells of connections from the network 12. The switch 10 comprises an output port mechanism 16 which sends cells of connections to the network 12. The switch 10 comprises a buffer mechanism 24 for storing cells. The buffer mechanism 24 is connected to the input port mechanism 14 and the output port mechanism 16. The switch 10 comprises a controller 18 which provides service to traffic of connections received at the input port mechanism 14 and which are to be sent out the output port mechanism 16. The controller 18 monitors the change in the number of cells in the buffer mechanism 24 and adjusts the service provided to the connections based on the change in the number of cells in the buffer mechanism 24. The controller 18 is connected to the buffer mechanism 24.

[0023] Preferably, the controller 18 also monitors the number of cells of connections received by the input port mechanism 14 and sent out the output port mechanism 16 and adjusting the service provided to the connections based also on the number of cells of connections received by the input port mechanism 14 and sent out the output port mechanism 16. The controller 18 preferably admits a new connection to the output port mechanism 16 of the switch 10 if

$$B'_{\text{eff}} + B'^{\text{int}}_{\text{eff}} + B'^{\text{act}}_{\text{act}} \leq B_{\text{link}}$$

where B'_{eff} is the effective bandwidth of the new connection, $B'^{\text{int}}_{\text{eff}}$ is the aggregate effective bandwidth of connections admitted within a current measurement interval int , $B'^{\text{act}}_{\text{act}}$ reflects an actual aggregate measured traffic on link 1 on the network 12 connected with the switch 10 determined at the beginning of the current interval, and B_{link} is the maximum

link capacity or line rate of 1. Preferably, the controller 18 updates B_{eff}^{int} during an interval if the connection B_{eff}^c is admitted by $B_{eff}^{int} = B_{eff}^{int} + B_{eff}^c$. B_{act} is preferably based on current link utilization, variation in link utilization, buffer occupancy, and rate of change of buffer occupancy.

[0024] The present invention pertains to a method for switching connections by a switch 10 of a network 12. The method comprises the steps of monitoring the traffic of connections received by the switch 10 and sent out the switch 10. Then there is the step of adjusting the service provided to the connections by the switch 10 based on the traffic of connections received by the switch 10 and sent out the switch 10.

[0025] Preferably, the monitoring step includes the step of monitoring the change in the number of cells in a buffer mechanism 24 of the switch 10 and adjusting the service provided to the connections based on the change in the number of cells in the buffer mechanism 24.

[0026] In the operation of the invention, a dynamic CAC is used to overcome the problem of under-utilization of link bandwidth. The scheme, using periodic measurements, keeps track of various utilization parameters. These include the current link utilization, the variation in link utilization, buffer occupancy, and rate of change of buffer occupancy. Based on the above measurements a dynamic estimate of the effective bandwidth of the aggregate traffic on link, denoted by B_{act}^l , is computed.

[0027] A new connection with effective bandwidth, B_{eff}^c , is admitted if:

$$B_{eff}^c + B_{act}^l \leq B_{link} \quad (2)$$

[0028] The above approach using measurements is more realistic and leads to link bandwidth being maximally allocated and utilized. To compare the performance of this approach with the previous approach, B_{eff}^l is tracked based on Eqn. 1. Thus, connections can be admitted in this scheme even with $B_{eff}^l > B_{link}$ for achieving better link bandwidth utilization. The flip side to this is the possible over-allocation of link bandwidth, leading to loss in QoS of connections.

[0029] Avoiding over-allocation in a measurement interval can be achieved in the following manner. Multiple new connections arriving between two measurements, based on Eqn. 2, will see the same last measured value of link utilization, B_{act}^l . Thus, the bandwidth taken up by admitted connections in a measurement interval is not reflected on subsequent admissions in the same interval. The above problem is solved by keeping track of the aggregate effective bandwidth of connections admitted within the interval, denoted by B_{eff}^{int} , and modifying Eqn. 2 as follows:

$$B_{eff}^c + B_{eff}^{int} + B_{act}^l \leq B_{link} \quad (3)$$

[0030] B_{eff}^{int} is reset to zero at the beginning of any interval, once the value of B_{act}^l has been updated based on measurement. During the interval if a connection is admitted, the parameter B_{eff}^{int} is updated as, $B_{eff}^{int} += B_{eff}^c$, otherwise it remains unchanged.

[0031] For a connection being torn down, a similar adjustment to B_{eff}^{int} can be made to reflect the freed up bandwidth. However, without actual per-connection usage statistics, it is not clear what the value of this credit should be. Clearly, B_{eff}^c may be an over-estimation, and should not be used. Thus, preferably, bandwidth is chosen not to be credited for connections torn down during an interval. An accurate impact of torn-down connections is obtained at the next link rate measurement.

[0032] Detecting and recovering from over-allocation in the long run can be accomplished. A link can also become over-allocated when some or all admitted connections send traffic corresponding to their B_{eff}^c . The system allows for this to be detected and then recovered from. Congestion due to over-allocation is detected based on two conditions occurring: a) the buffer occupancy for VBR/CBR traffic, denoted by Q_{vcbr} , crossing a certain high threshold, T , and b) B_{eff}^l , computed based on Eqn. 1 being greater than B_{link} . The condition is expressed as:

$$\text{if } (Q_{vcbr} > T) \text{ and } (B_{eff}^l > B_{link}) \text{ congestion = true; else congestion = false; (4)}$$

[0033] When over-allocation congestion is detected, recovery requires some connections to be released. A similar problem is also faced in an ATM network 12 using Inverse Multiplexing on ATM (IMA). In IMA, a link of larger capacity is emulated using multiple smaller capacity links. For example, a 6 Mbps link can be emulated using 4 DS-1 (1.5Mbps) links using reverse multiplexing. With one or more of these smaller capacity links going down, IMA requires the emulation to continue on the remaining active links. However, the remaining capacity may not be adequate to support all the connections. Thus, it also requires some connections to be dropped.

[0034] Currently, The ATM Forum Traffic Management Specification version 4.0, incorporated by reference herein, specifies traffic classes such as CBR, VBR, ABR, and, UBR. This classification determines the service received, in terms of rate, cell delay, and jitter, by a given connection in an ATM network 12. However, there is no notion of a *call release priority* based on which calls can be ordered while dropping to avoid congestion. Each connection in the system has such a *call release priority*. Such a priority could be either specified by the user at call set-up or determined by the

network 12 itself based on call characteristics.

[0035] When congestion is detected based on Eqn. 4, an estimate of the excess rate causing the congestion can be estimated from the rate of change of buffer occupancy and link utilization. Calls can then be released in decreasing order of their call release priority until congestion no longer holds. To minimize the number of calls being released, for a given priority, preferably the call with the closest match to the excess bandwidth is released first.

[0036] The details of a preferred embodiment of the algorithm is provided through the following pseudo-code.

Pseudo-code for a Task to Monitor Link Rates

10 [0037] The LINK_MONITOR_TASK task periodically monitors the transmitted CBR/VBR cell rate on each link. The period size is determined by a parameter LINK_MON_INTRVL_SIZE. The task keeps track of the measured rates from last N_INTRVLS (window size). Every interval a mean transmission rate and variance is computed based on a weighted function of the data points. For each link the task then uses the mean and variance to compute equivalent capacity.

```

15   LINK_MONITOR_TASK (LINK_MON_INTRVL_SIZE, N_INTRVLS, intrvl_weight [N_INTRVLS])
begin
  while (1)
20    curr_intrvl = (curr_intrvl+1) % N_INTRVLS

25    /* Grab semaphore for updating Link Capacities */
      semTake (linkMonitorSem);

30    for (linkid over all links)

35    /* Compute link cvbr tx rate in this intrvl */
      CellRate[linkid][curr_intrvl] = VCbrCellsSentThisIntrvl/LINK_MON_INTRVL_SIZE;

40    /* Compute mean rate based on data points from last N_INTRVLS */
      for (1 = 0; 1 < N_INTRVLS; 1++)
        win = (curr_intrvl + 1) % N_INTRVLS;
        mean += CellRate[linkid][win] * intrvl_weight[1];
45    end for

50    /* Compute variance of rate based on data points from last N_INTRVLS */
      for (1 = 0; 1 < N_INTRVLS; 1++)
        WIN = (CURR_INTRVL + 1) % N_INTRVLS;
        D = CellRate[linkid][win] - mean;
```

55

```

    var += D*D*intrvl_weight[1];

5   end for

    /* Update Link Capacity */

10  DCAC_LINK_EQUIV_CAPACITY (linkid, mean, var);

    end for

15  /* Release semaphore */
    semGive(linkMonitorSem);

20  /* Sleep until next measurement point */
    sleep (LINK_MON_INTRVL_SIZE);

25  end while

    end

30

```

Pseudocode for Dynamic Call Admission Algorithm

35 [0038] The CONNECTION_ADMISSION_CONTROL routine, for a new connection with given mean, peak, and burst_size computes its equivalent capacity. It then checks if the connection can be admitted without overallocating that link. It also updates the link capacity to reflect new call admitted in between measurements as discussed above.

40 [0039] To provide mutual exclusion in accessing the link capacity variables by the CONNECTION_ADMISSION_CONTROL and LINK_MONITOR_TASK task, semaphores are used. In the pseudocode below, a single semaphore for all links in the switch 10 is used. Other mutual exclusion mechanisms, with or without semaphores, can be envisioned. For example, even with semaphores, the granularity of mutual exclusion by associating a semaphore with each link capacity variable can be extended.

45

50

55

```
CONNECTION_ADMISSION_CONTROL (linkid, mean, peak, burst_size)
5 begin

10    /* Compute equivalent capacity of new connection */
        bufsz = reserved_buf_space_for_vbr [linkid];

15    CONNECTION_EQUIV_CAPACITY (mean, peak, burst_size, bufsz, ConnCapacity);

20    /* Grab semaphore for using and updating Link Capacity */
        semTake (linkMonitorSem);

25    /* Add the equiv capacity of the new connection to the equiv capacity of the link */
        DCACLinkCapacity [linkid] += ConnCapacity;

30    /* If link overallocated? revert change to link capacity and reject connection */
        bw = XLATE_EQUIV_CAP_TO_BANDWIDTH (LinkCapacity [linkid] );

35    if (bw > LinkLineRate)
        DCACLinkCapacity [linkid] -= ConnCapacity;

40

45

50

55
```

```

5           semGive (linkMonitorSem); /* Release semaphore */

10          return REJECT_CONNECTION;

/* Keep track of GCAC based link capacity for use in flow model based
10 * DCAC capacity estimation */
GCACLinkCapacity [linkid] += ConnCapacity;

15          /* Release semaphore */
semGive(linkMonitorSem);

20          return ACCEPT_CONNECTION;

end

25

```

Pseudocode for Equivalent Capacity Computations

- 30 [0040] In a preferred embodiment, two different equivalent capacity measures are kept track of. The first one based
on the *stationary model* maintains the mean and variance (M, V) of traffic. The second method is based on a *flow model*
maintains a parameter \hat{c} as discussed in R. Guerin, H. Ahmadi, and M. Naghshineh, incorporated by reference herein.
Equivalent capacity and its application to bandwidth allocation in high-speed networks 12. *IEEE Journal on Selected
Areas in Communication*, 9(7) :968-981, Sept. 1991, incorporated by reference herein. Both measures have the prop-
erty that they are additive. Thus, for the stationary model combining two different equivalent capacities, (M_1, V_1) and
(M_2, V_2) leads to $(M_1 + M_2, V_1 + V_2)$. Similarly, in the flow model combining two different equivalent capacities, (\hat{c}_1)
and (\hat{c}_2) leads to $(\hat{c}_1 + \hat{c}_2)$. The translation from the metric of equivalent capacity to required bandwidth (BW) is as fol-
lows. In the stationary model, $BW = M + \alpha\sqrt{V}$, where $\alpha = \sqrt{(-2\log(\varepsilon)\cdot\log(2\pi))}$, while in the flow model $BW = C$. The
value of BW used in an embodiment is minimum of the values derived from the two models.

40

45

50

55

```

/* Compute the equivalent capacity of a link based on its measured mean rate and variance */

5   DCAC_LINK_EQUIV_CAPACITY (linkid, mean, var)

begin
10    /* Update link capacity parameters for stationary mode, */
    DCACLinkCapacity[linkid].mean = mean;
        DCACLinkCapacity[linkid].var = var;

15    /* Update link capacity parameter for flow model */

20    t = GCACLinkCapacity[linkid].mean +  $\alpha\sqrt{GCACLinkCapacity[linkid].var}$ ;
        t1 = DCACLinkCapacity[linkid].mean +  $\alpha\sqrt{DCACLinkCapacity[linkid].var}$ ;
        DCACLinkCapacity[linkid].>flow_cap = GCACLinkCapacity[linkid].>flow_cap * (t1/t);

25    end

30    /* Compute the equivalent capacity of a new connection */

35    CONNECTION_EQUIV_CAPACITY (mean, peak, burst_size, bufsz, ConnCapacity)

40

45

50

55

```

```

begin

5      /* Update link capacity parameters for stationary model */
ConnCapacity.mean = mean;
10     ConnCapacity.var = var;

15     /* Update link capacity parameter for flow model */
ConnCapacity.flow_cap = c(peak, mean, burst_size, bufsz);

end

```

20

25 Example

[0041]

Modeling Assumptions

30

Traffic model

VBR entity parameters:

35 [0042]

Period between calls [0,100.00] ms
 Range of CAC_PCR [0.5,50.0] Mbps
 Range of CAC_SCR [0.45,CAC_PCR] Mbps
 40 Range of CAC_MBS [1 cell, 1000] us
 Mean CAC_MIS = mean CAC_MBS
 Range of ACT_PCR [0.45,CAC_PCR] Mbps
 Range of ACT_MBS [5, CAC_MBS] us
 Range of ACT_MIS [CAC_MIS, 250+CAC_MIS] us
 45 Range of transfer size [0.1,100] Mbits

Link Parameters for equiv_cap computation

[0043]

50

VBR: queue sz 256, VBR_OB 100, VBR_BUFF_OB 100

Link Parameters for DCAC

55 [0044]

Max. allowed util by VBR 0.80 LR
 Link monitoring interval 5.000000 ms

Monitoring data maintained for 4 windows: Window weights: 0.50 0.25 0.15 0.10

[0045] Figures 1a-1g show characteristics of calls generated in simulations. Plots depict distribution of pcr, scr, and mbs requested by calls during CAC. The distribution of actual pcr, scr, and mbs generated by calls are also depicted alongside. The distribution of the transfer size is also depicted. In this experiment, total number of calls generated were ~ 45,000. Note, these statistics are for all calls, accepted/rejected.

[0046] Figures 2a-2e show plots which depict observed number of calls admitted by gcac, number of calls rejected, gcac parameters, link utilization, and vbr buffer occupancy. Note average link utilization with GCAC is around 150,000 cells/s or 40% line rate.

10

Experiment 1: GCAC

15

[0047] Figures 3a-3f depicts the plots for observed number of calls admitted by GCAC, number of calls rejected, gcac parameters, link utilization, and vbr buffer occupancy. Note average link utilization with GCAC is around 150,000 cells/s or 40% line rate. From these observations, it can be concluded that for the calls being generated in this experiment, the average value of CAC was only 40% of requested value.

Experiment 2: DCAC

20

[0048] The above experiment was repeated with DCAC. The link bandwidth available to DCAC was set to 80% of line rate. The idea being to reserve some bandwidth, in this case 20% line rate, to allow for admitted calls to intermittently burst up to their requested CAC parameters.

25

[0049] Although the invention has been described in detail in the foregoing embodiments for the purpose of illustration, it is to be understood that such detail is solely for that purpose and that variations can be made therein by those skilled in the art without departing from the spirit and scope of the invention except as it may be described by the following claims.

Claims

30

1. A switch for a network comprising:

an input port mechanism which receives traffic of connections from the network, each connection having a priority;

35

an output port mechanism which sends traffic of connections to the network; and

40

a controller which serves connections and which monitors the connections received by the input port mechanism and sent by the output port mechanism and releases connections according to a connection's priority when a predetermined condition exists in the switch, said controller connected to the input port mechanism and the output port mechanism, each connection requesting a specific bandwidth from the controller.

2. A switch as described in Claim 1 wherein the predetermined condition includes an overload condition of the switch.

45

3. A switch as described in Claim 2 wherein the overload condition includes receiving more traffic on connections received by the input port mechanism than the controller can process at a given time.

4. A switch as described in Claim 3 wherein the controller determines excess bandwidth associated with the overload condition, and the controller releases the connection, for a given priority, with the closest match to the excess bandwidth.

50

5. A switch as described in Claim 4 wherein the controller releases connections according to the connections' traffic class, where UBR connections are released first, then ABR connections are released next, then VBR connections are released next and then CBR connections are released from the switch.

55

6. A switch as described in Claim 5 wherein the input port mechanism comprises I input ports connected to the controller, where I is an integer and greater than or equal to 1; and wherein the output mechanism comprises O output ports connected to the controller, where O is an integer greater than or equal to 1.

7. A switch as described in Claim 6 wherein the controller releases the connection with the lowest priority.
8. A switch as described in Claim 7 wherein the controller releases connections whose sum of requested bandwidths most closely matches the excess bandwidth.
9. A switch as described in Claim 8 including a buffer mechanism connected to the output port mechanism and wherein the overload condition exists if the sum of requested bandwidths of connections on an output port is greater than the output port's bandwidth and buffer resources associated with it are being exceeded.

10 if ($Q_{vcbr} > T$) and ($B'_{eff} > B_{link}$)

where Q_{vcbr} is the buffer occupancy for VBR/CBR traffic, T is a threshold, B'_{eff} is the aggregate effective bandwidth of connections on link 1, and B_{link} is the maximum link capacity or line rate.

- 15 10. A method for switching connections comprising:

monitoring traffic of connections received by a switch;

20 releasing connections from the switch according to the connection's priority when a predetermined condition in the switch exists.

11. A method as described in Claim 10 wherein the monitoring step includes the step of monitoring traffic of connections received by and sent out from the switch.

- 25 12. A method as described in Claim 11 wherein the releasing step includes the step of releasing connections when the predetermined condition of overload exists in the switch where there is more traffic of connections received by the switch than the switch can process.

- 30 13. A method as described in Claim 12 wherein the releasing step includes the steps of determining, when the overload condition exists, the excess bandwidth associated with the overload condition; choosing a connection, for a given priority, with the closest match to the excess bandwidth; and releasing the connection with the closest match to the excess bandwidth.

- 35 14. A method as described in Claim 13 wherein the choosing step includes the step of choosing the connection with the lowest priority.

15. A method as described in Claim 14 wherein the choosing step includes the step of choosing the connections whose sum of requested bandwidths most closely matches the excess bandwidth.

- 40 16. A method as described in Claim 15 wherein the choosing step includes the step of choosing a connection according to traffic class for release, where UBR connections are released first, then ABR connections are released next, then VBR connections are released next and then CBR connections are released.

- 45 17. A switch for a network comprising:

an input port mechanism which receives traffic of connections from the network;

an output port mechanism which sends traffic of connections to the network; and

- 50 a controller which provides service to traffic of connections received at the input port mechanism and which are to be sent out the output port mechanism, said controller monitoring the traffic of connections received by the input port mechanism and sent out the output port mechanism and adjusting the service provided to the connections based on the traffic of connections received by the input port mechanism and sent out the output port mechanism, said controller connected to the input port mechanism and the output port mechanism.

- 55 18. A switch for a network comprising:

an input port mechanism which receives cells of connections from the network;

an output port mechanism which sends cells of connections to the network;

a buffer mechanism for storing cells, said buffer mechanism connected to the input port mechanism and the output port mechanism; and

a controller which provides service to traffic of connections received at the input port mechanism and which are to be sent out the output port mechanism, said controller monitoring the change in the number of cells in the buffer mechanism and adjusting the service provided to the connections based on the change in the number of cells in the buffer mechanism, said controller connected to the buffer mechanism.

10 19. A switch as described in Claim 18 wherein the controller also monitors the number of cells of connections received by the input port mechanism and sent out the output port mechanism and adjusting the service provided to the connections based also on the number of cells of connections received by the input port mechanism and sent out the output port mechanism.

15 20. A switch as described in Claim 19 wherein the controller admits a new connection to the output port mechanism of the switch if

$$B^c_{\text{eff}} + B^{int}_{\text{eff}} + B^l_{\text{act}} \leq B_{\text{link}}$$

20 where B^c_{eff} is the effective bandwidth of the new connection, B^{int}_{eff} is the aggregate effective bandwidth of connections admitted within a current measurement interval int , B^l_{act} reflects an actual aggregate measured traffic on link I on the network connected with the switch determined at the beginning of the current interval, and B_{link} is the maximum link capacity or line rate of 1.

25 21. A switch as described in Claim 20 wherein the controller updates B^{int}_{eff} during an interval if the connection B^c_{eff} is admitted by $B^{int}_{\text{eff}} = B^{int}_{\text{eff}} + B^c_{\text{eff}}$.

30 22. A switch as described in Claim 21 wherein B_{act} is based on current link utilization, variation in link utilization, buffer occupancy, and rate of change of buffer occupancy.

23. A method for switching connections by a switch of a network comprising the steps of:

35 monitoring the traffic of connections received by the switch and sent out the switch; and

36 adjusting the service provided to the connections by the switch based on the traffic of connections received by the switch and sent out the switch.

40 24. A method as described in Claim 23 wherein the monitoring step includes the step of monitoring the change in the number of cells in a buffer mechanism of the switch and adjusting the service provided to the connections based on the change in the number of cells in the buffer mechanism.

45

50

55

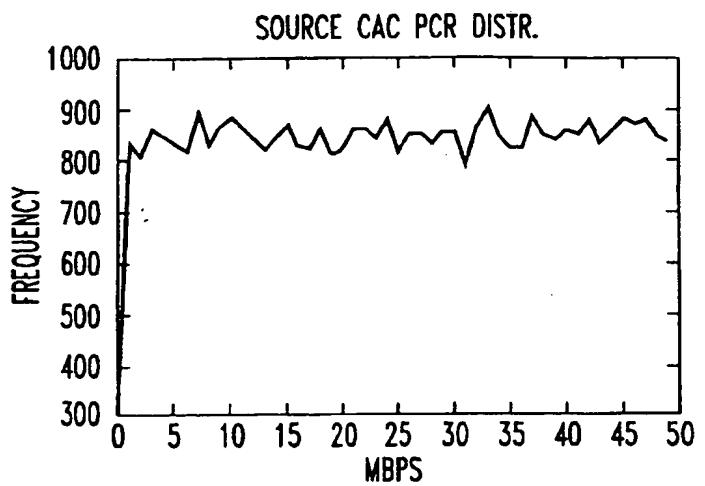


FIG.1a

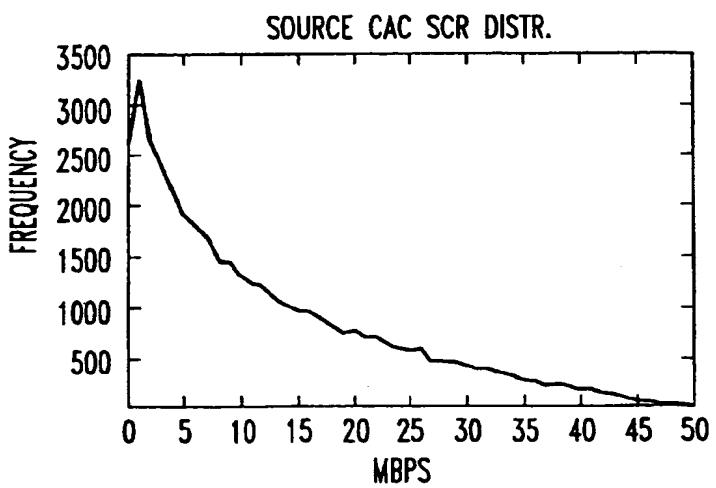


FIG.1b

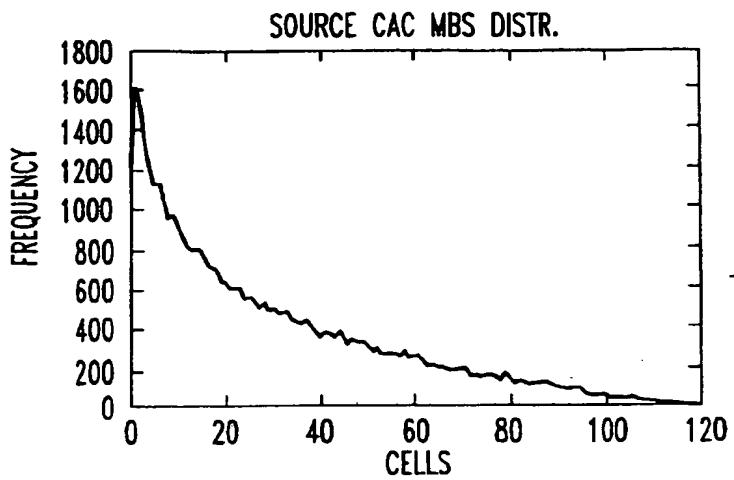


FIG.1c

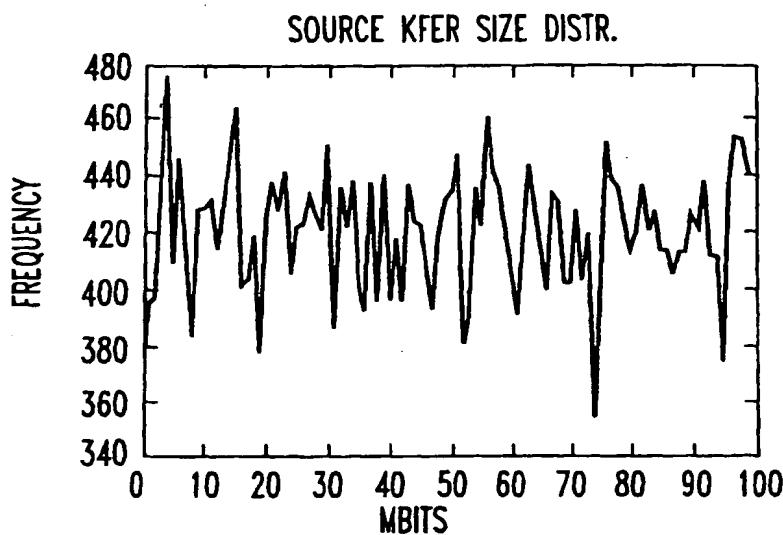


FIG.1d

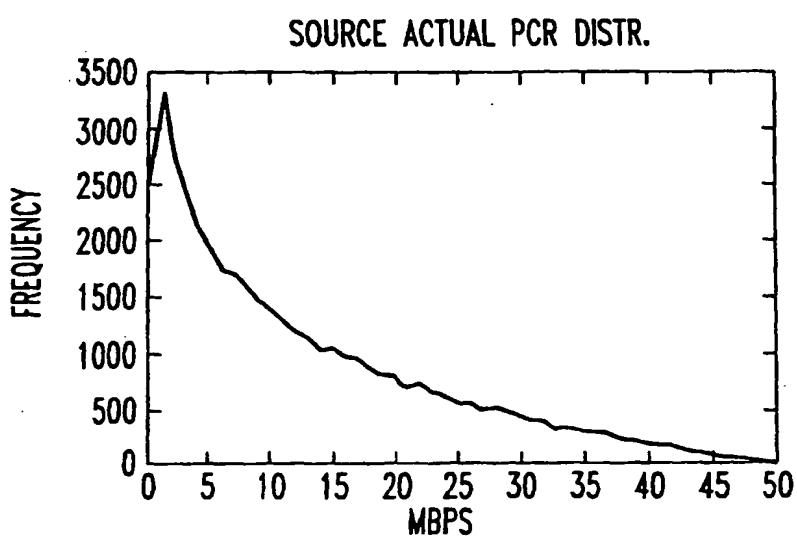


FIG.1e

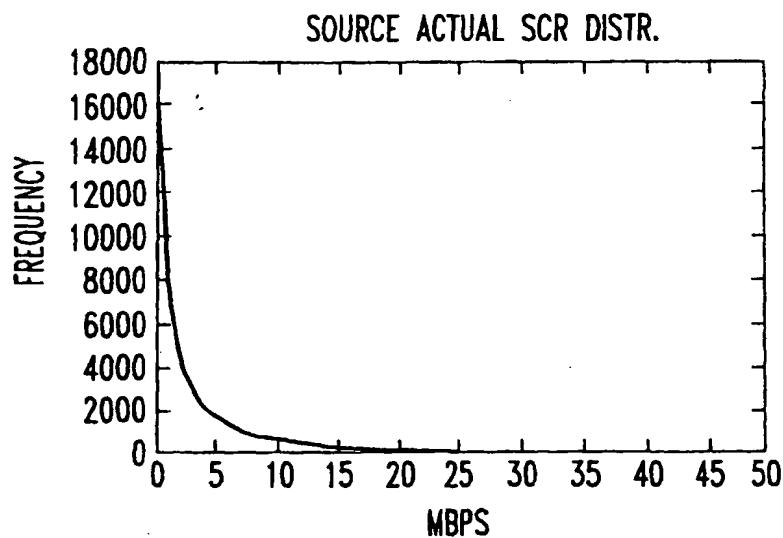


FIG.1f

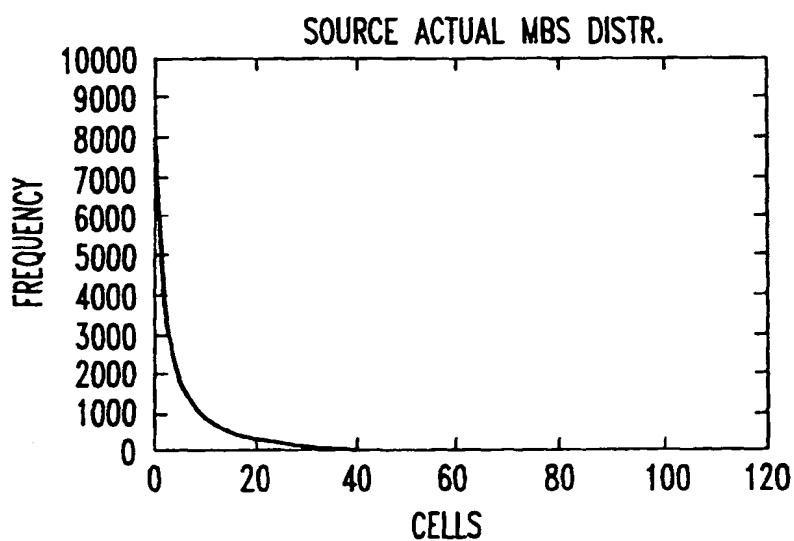


FIG.1g

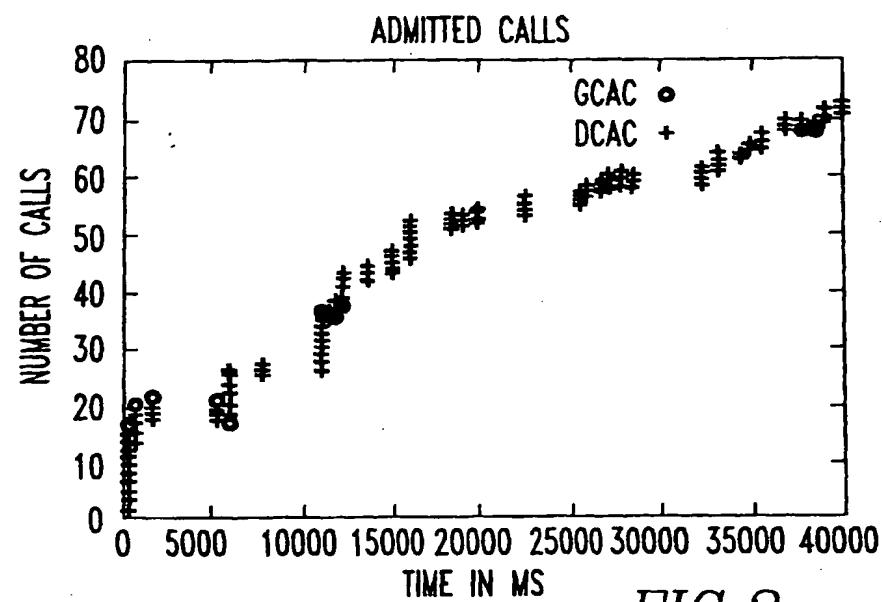


FIG.2a

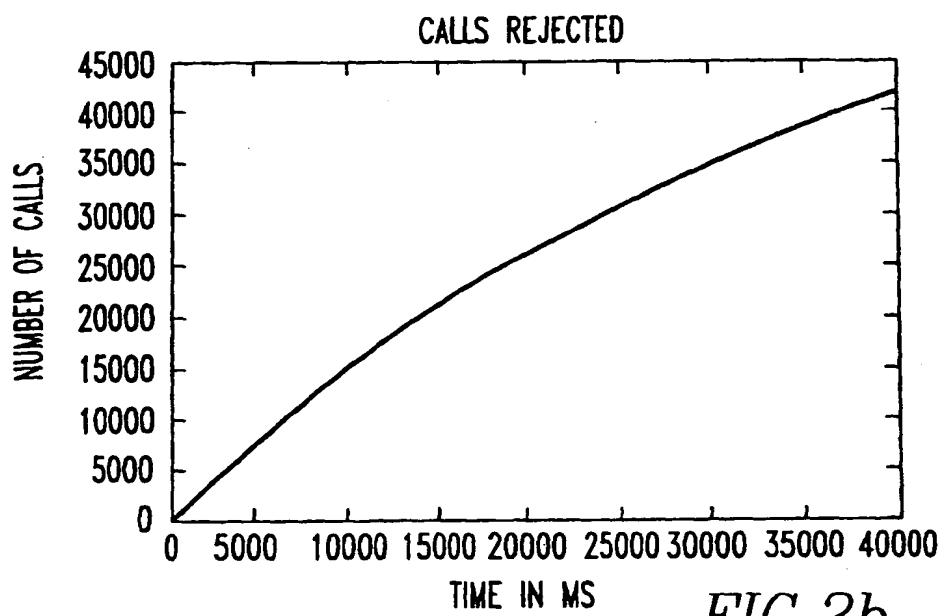


FIG.2b

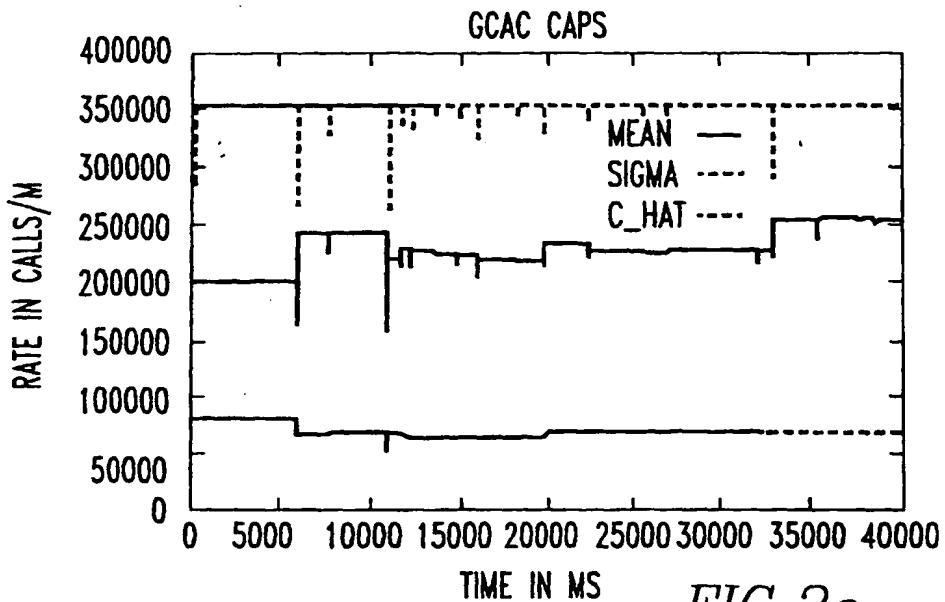


FIG.2c

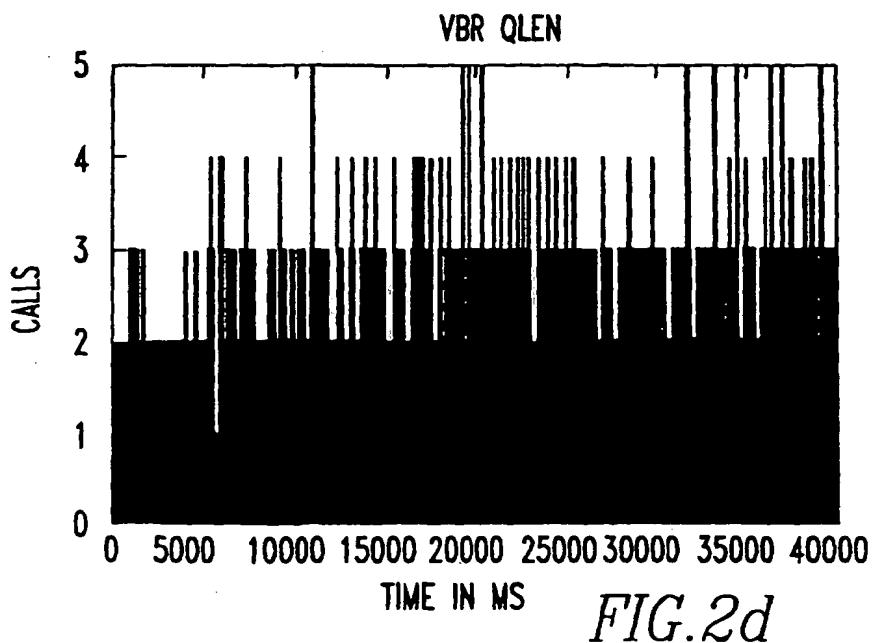
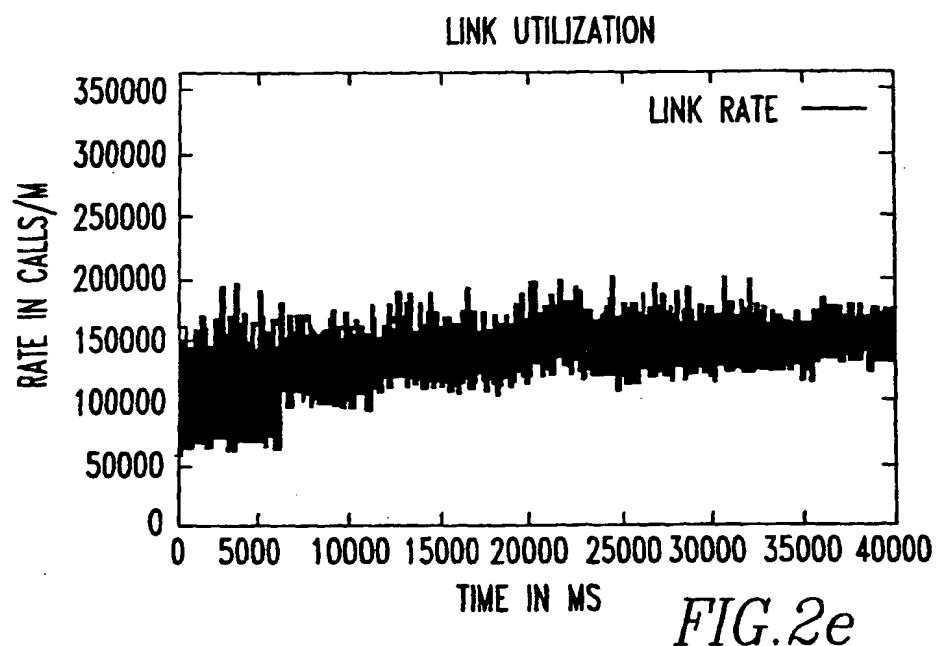


FIG.2d



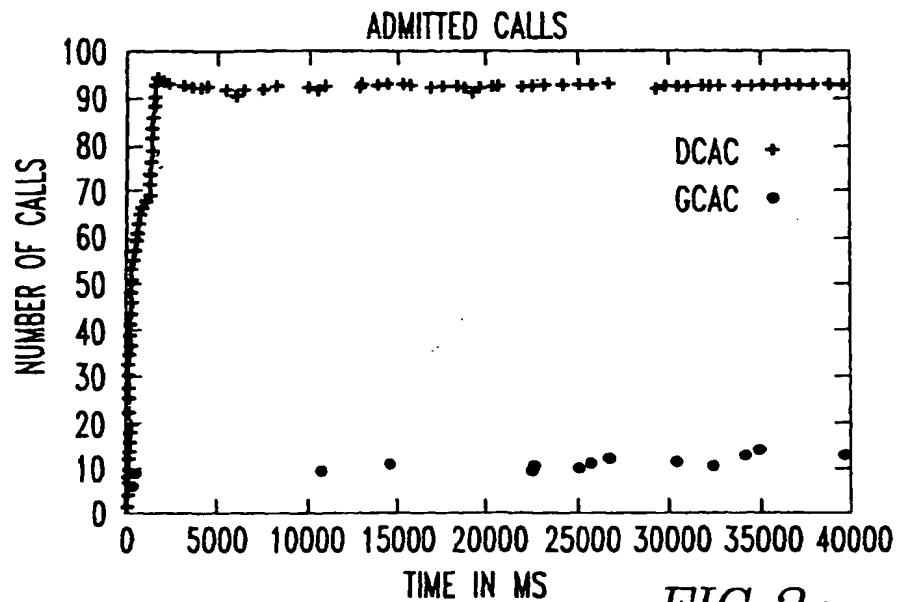


FIG. 3a

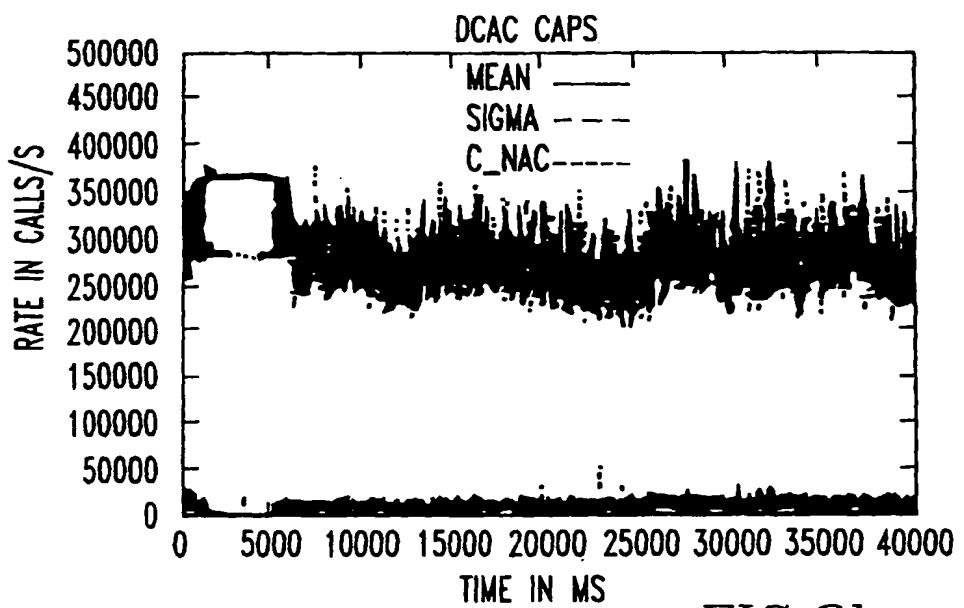
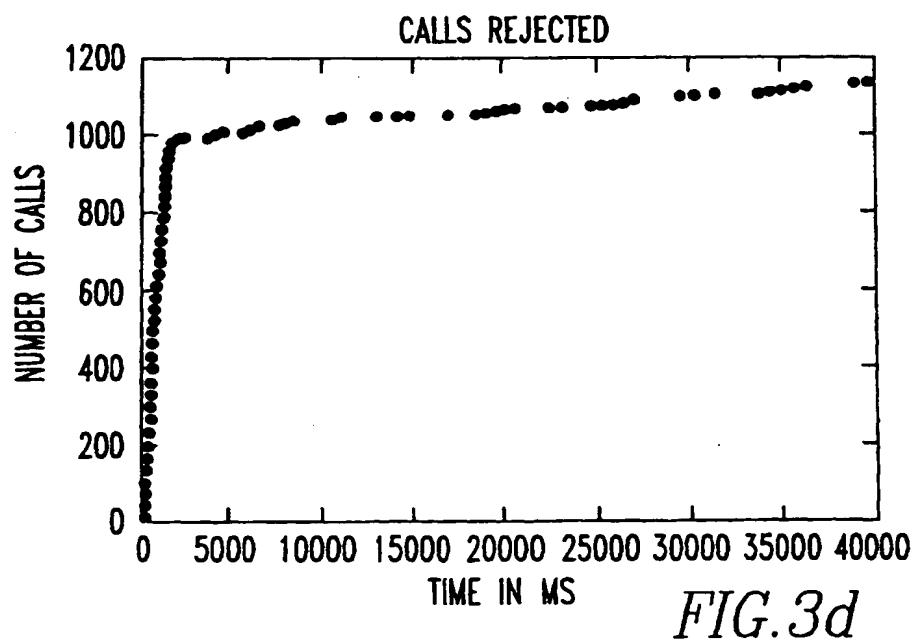
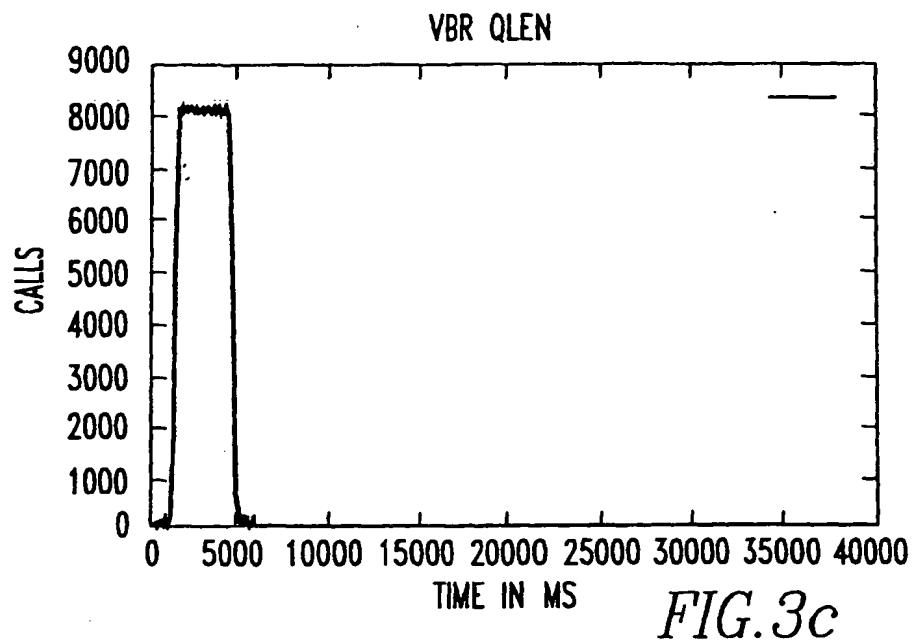


FIG. 3b



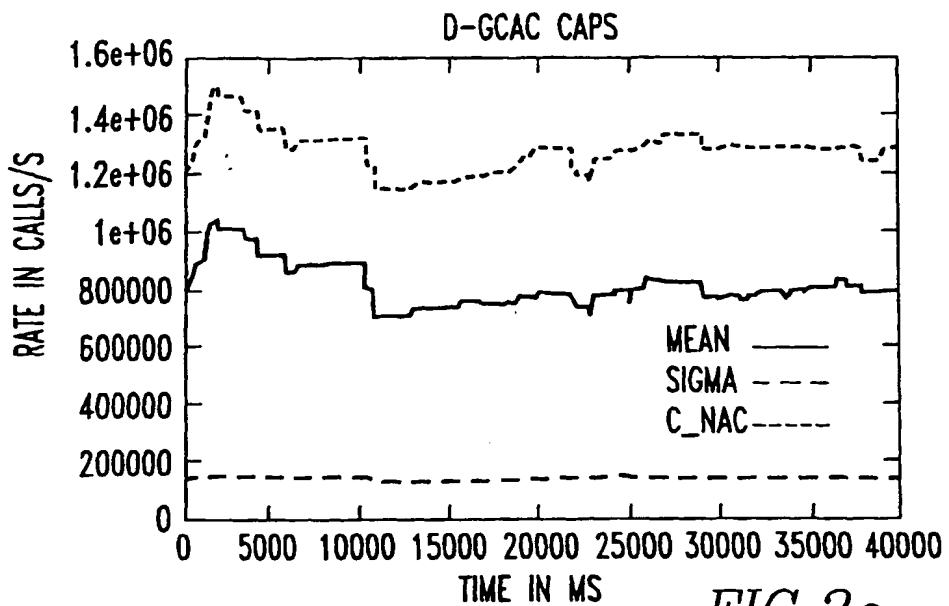


FIG. 3e

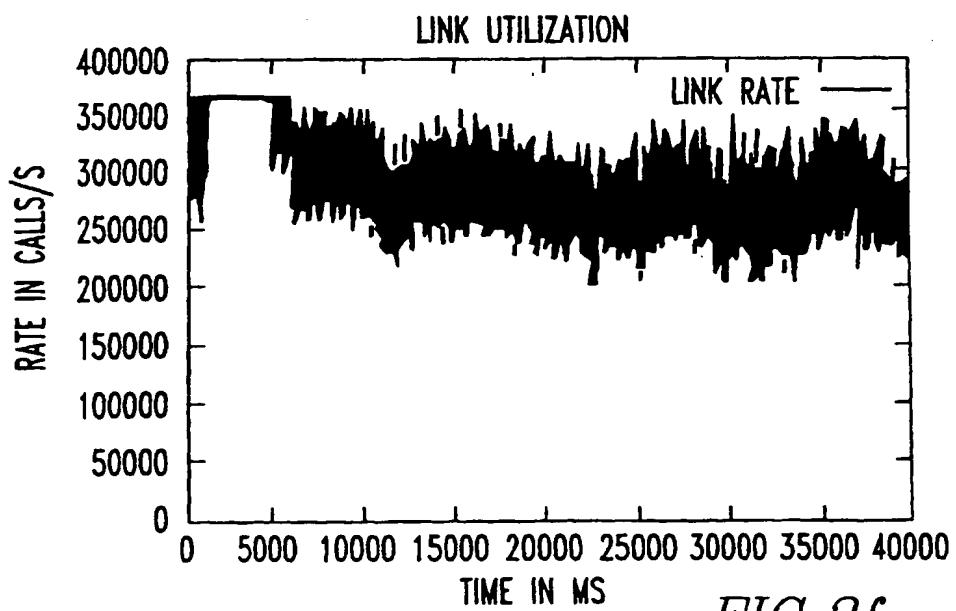


FIG. 3f

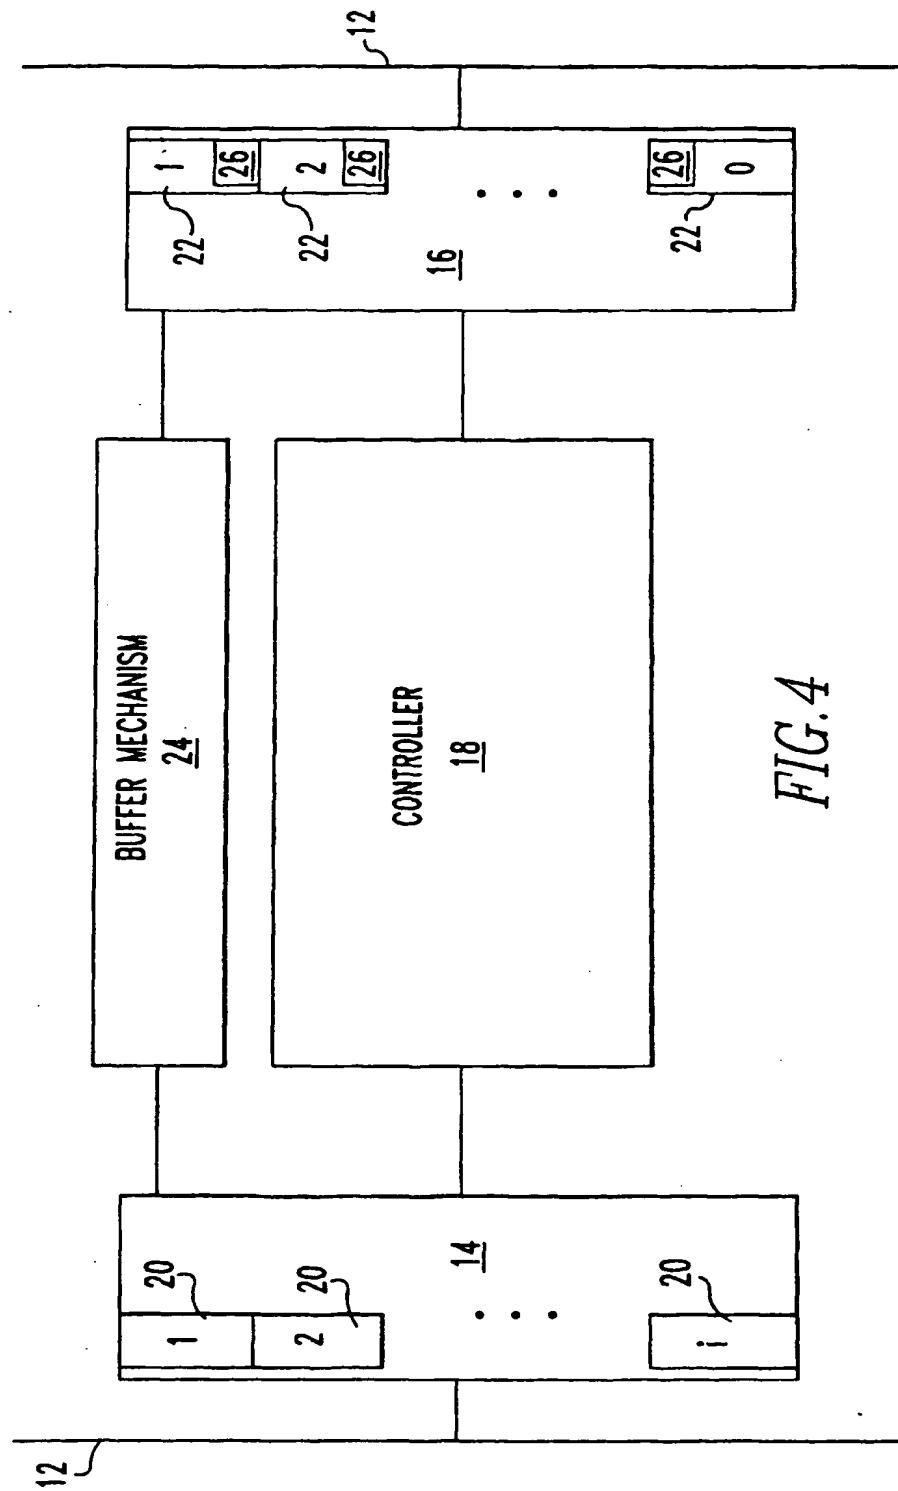


FIG. 4